

# Exemplar-Free Continual Transformer with Convolutions

Anurag Roy<sup>1</sup> Vinay K. Verma<sup>2,†</sup> Sravan Vonna<sup>1</sup> Kripabandhu Ghosh<sup>3</sup> Saptarshi Ghosh<sup>1</sup> Abir Das<sup>1</sup>  
<sup>1</sup>IIT Kharagpur, <sup>2</sup>IML Amazon India, <sup>3</sup>IISER Kolkata

{anurag\_roy@,vg.sravan@,saptarshi@cse.,abir@cse.}@iitkgp.ac.in, {vinayugc,kripa.ghosh}@gmail.com

## Abstract

Continual Learning (CL) involves training a machine learning model in a sequential manner to learn new information while retaining previously learned tasks without the presence of previous training data. Although there has been significant interest in CL, most recent CL approaches in computer vision have focused on convolutional architectures only. However, with the recent success of vision transformers, there is a need to explore their potential for CL. Although there have been some recent CL approaches for vision transformers, they either store training instances of previous tasks or require a task identifier during test time, which can be limiting. This paper proposes a new exemplar-free approach for class/task incremental learning called ConTraCon, which does not require task-id to be explicitly present during inference and avoids the need for storing previous training instances. The proposed approach leverages the transformer architecture and involves re-weighting the key, query, and value weights of the multi-head self-attention layers of a transformer trained on a similar task. The re-weighting is done using convolution, which enables the approach to maintain low parameter requirements per task. Additionally, an image augmentation-based entropic task identification approach is used to predict tasks without requiring task-ids during inference. Experiments on four benchmark datasets demonstrate that the proposed approach outperforms several competitive approaches while requiring fewer parameters.<sup>1</sup>

## 1. Introduction

Humans excel at solving newer tasks without forgetting previous knowledge. Deep neural networks, however, face the challenge of forgetting old knowledge when trained for a novel task. This problem, known as *Catastrophic Forgetting*, is a fundamental challenge if a model needs to learn when data arrives in a sequence of non-overlapping tasks. Con-

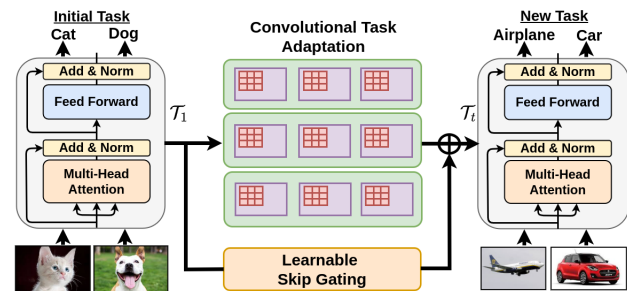


Figure 1: The proposed ConTraCon architecture.  $T_1$  represents a transformer model trained on a previous task which is adapted to a new task using only a few learnable parameters and temporal skip-gating to a new transformer  $T_t$  on the Right.

tinual Learning [1, 21, 26, 37, 40, 41] aims to handle such problems arising from non-stationary data to retain previously learned knowledge as well as acquire knowledge from new data with limited or no access to previous data.

Although continual learning in computer vision has witnessed remarkable progress, most of the methods are tailored for CNNs. Recently, vision transformers have shown promising results in big data regime [8, 15, 32]. However, data hungry vision transformers and data-scarce continual learning do not seem to go hand in hand, and thus continual learning in vision transformers has received relatively little attention. Early continual learning approaches on ConvNets relied on exemplar rehearsal which re-trains newer models on previous data instances stored in a fixed size memory buffer [2, 7, 9, 23, 33, 38, 40, 41]. Dytox [16] and LVT [52] are contemporary works addressing continual learning on vision transformers using such previously stored data. However, storing task samples in raw format may not always be feasible, especially for tasks where long-term storage of data is not permitted owing to privacy or data use legislation [24]. *Dynamic architecture* based methods, on the other hand, grows an initial model dynamically or rearranges its internal structure, [19, 44, 46, 58] with the arrival of new tasks without the need for storing previous data.

The main focus of this paper is in modeling the new set of parameters for new tasks with as low overhead as possible.

<sup>†</sup>Work started before joining Amazon

<sup>1</sup>Project Page: <https://cvir.github.io/projects/contracon>

We conjecture that better representation learning capability of a transformer keeps it ready for easily adapting to new tasks by *manipulating* the already learned weights. Convolution offers a cheap way to manipulate transformer weights and the weight matrices are natural fits to convolution as input and output. In this paper, we propose a dynamically expandable architecture with novel convolutional adaption on previously learnt transformer weights to obtain new weights for new tasks. In addition to this, we employ a learnable skip-gating which learns to convexly combine the old and the convoluted weights. This helps significantly in maintaining the stability-plasticity tradeoff by balancing between how much to retain and how much to forget. The resulting proposed approach – **Continual Transformer with Convolutions** (ConTraCon) (ref. Fig. 1) – not only leverages a vision transformer’s learning capability but also does it with significantly low memory overhead per task as is shown in our experiments.

Specifically, for each new task, we reweigh the key, query, and value weights of the previously learned transformer by convolving them with small filters. Transformers are known for their ability to capture long-range dependencies between patches. Convolution, on the other hand, exploits a local neighborhood only. Such local dependency not only restricts large changes in the weights for the new task from the old tasks, but also allows us to achieve this with very little increase in the model size. As the convolution weights are separate for different tasks, during inference, this will require the information of which task an image belongs to. However, in class incremental continual learning, images may come arbitrarily without associated task information. To tackle the challenging scenario, we propose a novel entropy-based criterion to infer the task before calling for the corresponding task-specific convolution weights. In particular, our proposed approach creates multiple augmented views of a test image and evaluates the agreement of their predictions among different task-specific models. The task-id of the image is determined by identifying the task giving the lowest entropy value of the average predictions from the various augmented views. The task-specific model with the most consistent and confident predictions across different augmentations corresponds to the correct task for the image.

Extensive experiments on four benchmark datasets, considering both the availability and unavailability of task-ids at test time, demonstrate the superiority of our proposed method over several state-of-the-art methods including the extension of popular approaches on CNNs to transformers. Despite being exemplar-free, our approach outperforms state-of-the-art exemplar-based continual learning approaches that use transformers as backbone architectures (Dytox [16] and LVT [52]) with an average improvement of 5%. Moreover, our approach accomplishes this with only about  $\sim 60\%$  parameters used by above models. We also perform ablation

studies highlighting the contribution of each component in our approach. To summarize, our key contributions include:

- We propose ConTraCon, a dynamic architecture for CL on transformers. We use convolution on the Multi-head self-attention (MHSA) layers to learn new tasks thereby achieving better performance with significantly low memory overhead. We also apply a temporal skip-gating function to tradeoff between stability and plasticity.
- Our entropy based approach adds the flexibility of not having to know the task information during inference.
- We performed extensive experimentation and ablation of the proposed approach thereby validating the superiority of our model and helping to disentangle the significance of different components.

## 2. Related Work

**Continual Learning:** CL approaches can be broadly classified into – 1) Exemplar-replay methods, 2) regularization methods and 3) dynamic architecture methods. To avoid forgetting when learning a new task, replay approaches repeat past task samples that are kept in raw format [2, 4, 6, 7, 9, 16, 25, 40, 52] or generated with a generative model [45]. Usually, replay-based approaches have a fixed memory which stores samples. Regularization-based approaches, on the other hand, prevent catastrophic forgetting by encouraging important parameters to lie in close vicinity of previous solutions with the introduction of penalty terms to the loss function [26, 31, 59] or constraining the direction of parameter update [18, 43]. However, the regularization prevents the model from performing well for long task sequences. Dynamic Architecture methods [51, 57, 58] grow the model dynamically with the arrival of new tasks. Although these approaches can learn long sequences of tasks, these can suffer from the huge memory and compute overhead if not managed properly.

**Transformers:** Transformers [50] introduced multi-head self-attention (MHSA) in the form of encoder-decoder architecture for machine translation and have since become the state-of-the-art in many NLP tasks [3, 5, 10, 13, 14]. In computer vision, the pioneering Vision Transformer (ViT) [15] directly applied a transformer encoder to image classification with image patches as input. There have been multiple works since then, including DeiT [47], ConViT [11] CCT [22] *etc.* with architecture and training procedures modifications.

**Continual Learning in Vision Transformers:** Despite the success of vision transformers, continual learning with vision transformers has received very little attention. Recently Dytox [16] proposed to learn new tasks through the expansion of special tokens known as task tokens. Another recent approach, LVT [52], proposed an inter-task attention mechanism that absorbs the previous tasks’ information and slows down the drift of information between previous and current tasks. Both Dytox and LVT require extra memory for stor-

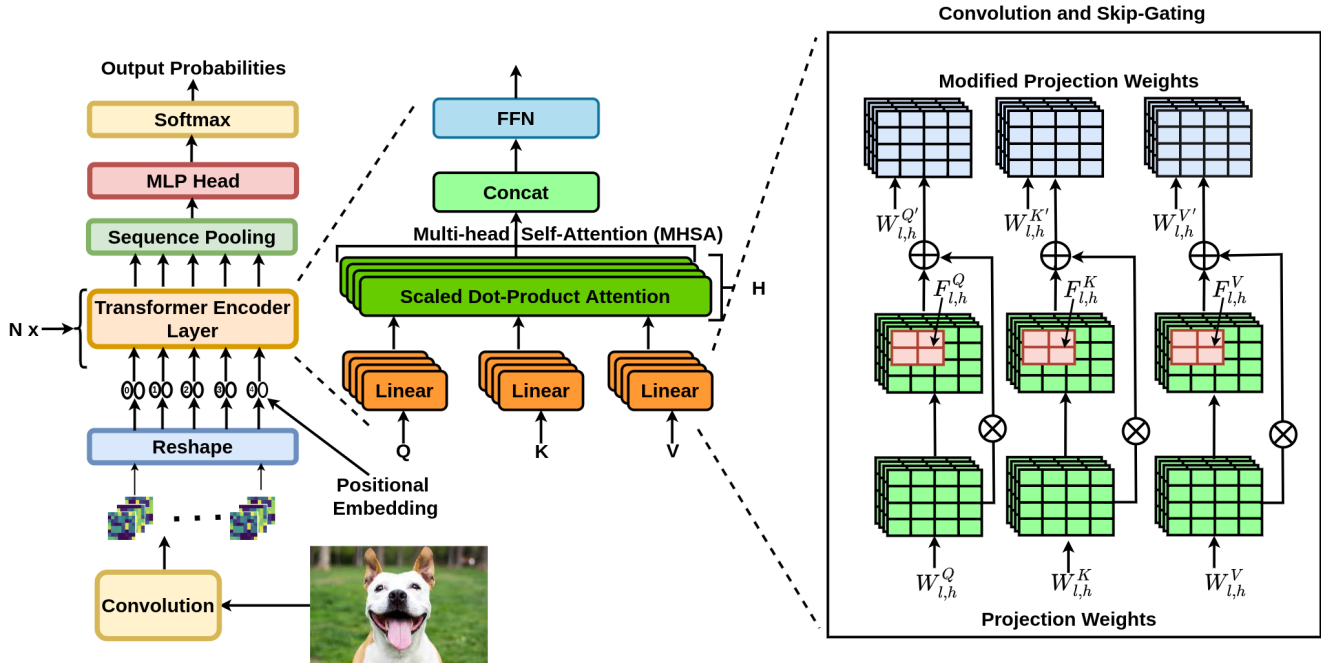


Figure 2: Illustration of the proposed ConTraCon model. **Left** part contains the pictorial representation of the CCT model. **Right** part represents the Convolution and skip-gating operation used to adapt a pre-trained CCT model to new tasks. To learn a new task, the existing set of projection weights (highlighted in green) are modified by means of adaptable task-specific convolution filters (highlighted in pink) and learnable skip-gate (denoted by  $\otimes$ ), resulting in a fresh collection of projection weights (highlighted in blue). This eliminates the necessity of developing a new set of projection weights for each new task, and considerably decreases the number of task-specific parameters.

ing training instances from previous tasks. Another recent method, MEAT [56] uses learnable masks to help isolate previous tasks’ parameters that are important for learning new tasks. However, as the architecture is not expandable, the number of tasks the model can learn is limited. Additionally, task-ids are required during inference. A few recent approaches learn soft prompts on new tasks rather than directly modifying the vision encoder which is pre-trained on large amounts of training data. L2P [54] and DualPrompt [53] learn a query mechanism to select task relevant prompts from a pool of prompts structured in a key-value scheme. One drawback of these approaches is they require large pre-training on huge datasets and may fail if the pre-training data is not representative of the later tasks in hand.

Different from the contemporary works, ConTraCon uses convolution in an intelligent way on the weights that does not require a rehearsal memory for storing samples. In addition, our proposed approach is flexible enough to operate without the knowledge of task-ids by comparing the entropy of average predictions of several augmented views. It does not depend on large pre-training and can gracefully handle continual learning scenarios with largely varying tasks.

### 3. Method

In this section, we first describe our problem setting, and then describe the proposed ConTraCon model.

### 3.1. Problem Setting

Considering  $T$  tasks arriving sequentially, the goal is to obtain a model which is capable of learning new tasks without forgetting the previous ones. Specifically, for a task  $t \in \{1, 2, \dots, T\}$  the model is exposed to training samples  $\{x_i^t, y_i^t\}$ , where  $x_i^t$  is the  $i^{\text{th}}$  sample of the  $t^{\text{th}}$  task and  $y_i^t \in C^t$  is the corresponding label belonging to the set of classes  $C^t$  with the set of the class labels being mutually exclusive, i.e.,  $C^0 \cap C^1 \dots \cap C^T = \phi$ . Our approach is flexible to work in two common CL scenarios, namely (1) Task Incremental Learning (TIL), and (2) Class Incremental Learning (CIL). In TIL, task-ids for each input is assumed to be known at test time. In the more challenging CIL scenario, task-ids are unknown and need to be inferred [12, 49]. Our approach is also *exemplar-free*, i.e., data from previous tasks is *not* available for any subsequent use.

### 3.2. Preliminaries

The vision transformer (ViT) [15] has three major components: (1) tokenizer layer, (2) multi-head self-attention (MHSA) and (3) feed-forward network (FFN). In ViT, an image is divided into  $n$  fixed-size non-overlapping patches. Each of these patches are embedded into  $d$ -dimensional vectors. Formally, the input image  $x_i^t$  is tokenized into  $n$  patches  $z_0 \in \mathbb{R}^{n \times d}$ . Positional embeddings are used to add spatial information in the token sequence. In ViT, a learnable class

embedding is used for the final classification.

A single encoder layer of ViT consists of stacks of MHSA, layer normalization, and FFN blocks, each with residual connections. At the  $l^{th}$  layer, the input is  $z_l$  generating the output  $z_{l+1}$  that goes into the next layer as input,  $l \in \{1, 2, \dots, L\}$  where  $L$  is the total number of encoder layers. Each MHSA block consists of  $H$  separate self-attention heads. At the  $l^{th}$  layer, self-attention values from head  $h \in \{1, 2, \dots, H\}$  is:

$$A_{l,h}(Q_{l,h}, K_{l,h}, V_{l,h}) = \text{softmax}\left(\frac{Q_{l,h}K_{l,h}^T}{\sqrt{d_k}}\right)V_{l,h} \quad (1)$$

where  $Q_{l,h} = z_l W_{l,h}^Q$ ,  $K_{l,h} = z_l W_{l,h}^K$  and  $V_{l,h} = z_l W_{l,h}^V$  are query, key, and value with learnable weights  $W_{l,h}^Q$ ,  $W_{l,h}^K$  and  $W_{l,h}^V \in \mathbb{R}^{d \times d_k}$  respectively with  $d_k$  being the dimension of key, query and value vectors.  $A_{l,h}$  is the attention matrix obtained as dot product between query and key. The activations from different attention heads are then concatenated and passed through a linear projection layer as:

$$MHSA_l(Q, K, V) = \text{Concat}(A_{l,1}, A_{l,2}, \dots, A_{l,H})W^O \quad (2)$$

The FFN block consists of two linear layers and an activation function (usually GELU).

### 3.3. ConTraCon

We now describe the different modules of ConTraCon. An overview of the approach is illustrated in Fig. 2.

**Convolution on Transformer Weights:** Self-attention in transformers plays a significant role in their improved performance. The model learns the relationship between all image patches. The large effective receptive field of the self-attention layers of transformers helps them to learn the visual representations required for vision tasks [39] by exploiting large amounts of labeled training data. However, in the presence of only a handful of labeled training data, one of the requirements for capturing the new knowledge without overfitting is to employ only a small number of learnable parameters. We propose task-specific convolution operations to reweigh the previously learned key, query, and value weights. The sharable computation offered by the convolutional filters prevents overfitting and maintains plasticity while learning new tasks. On the other hand, the local computation offered by them helps to maintain the required stability by restricting large changes in the weights. Additionally, convolutional filters being small, help to keep the memory footprint of the newly learned models in check as well.

Specifically, on a ViT, with  $L$  encoders and  $H$  attention heads per MHSA block, each new task is continually learned by applying convolution kernels on each of the self-attention

layers of the MHSA blocks in the transformer:

$$\begin{aligned} W_{l,h}^{Q'} &= \text{Conv}(F_{l,h}^Q, W_{l,h}^Q) \\ W_{l,h}^{K'} &= \text{Conv}(F_{l,h}^K, W_{l,h}^K) \\ W_{l,h}^{V'} &= \text{Conv}(F_{l,h}^V, W_{l,h}^V) \end{aligned} \quad (3)$$

where  $\text{Conv}(\cdot)$  denotes the convolution operation which is applied on the weights  $W_{l,h}^Q$ ,  $W_{l,h}^K$  and  $W_{l,h}^V$  with filters  $F_{l,h}^Q$ ,  $F_{l,h}^K$  and  $F_{l,h}^V$  respectively. Note that the filter weights are learned separately for each task  $t \in \{1, 2, \dots, T\}$ . However, to avoid clutter, we are not incorporating the task index above. These cheap operations help the transformers learn a similar task without forgetting previous tasks and simultaneously keep the compute and memory footprint low.

**Learnable Skip-Gate:** It has been shown in ResCL [30] that a residual connection combining previous and newly learned weights could improve continual learning performance. Taking inspiration, we took a step further and replaced the heuristic combination with a learnable skip connection where the combined weight is learned directly from data. Specifically, we apply the skip connections as follows.

$$W_{l,h}^{Q''} = W_{l,h}^{Q'} + \sigma(\alpha_{l,h}) * W_{l,h}^Q \quad (4)$$

Where  $\alpha_{l,h} \in \mathbb{R}$  is a learnable parameter and  $\sigma(\cdot)$  is the sigmoid function.  $W_{l,h}^{Q''}$  is the weight matrix of the query for layer  $l$  and head  $h$  corresponding to the new task. Here also, we drop the task index  $t$  to avoid clutter. A similar operation is performed on the key and values weight matrices.

### 3.4. Task Prediction

In the class incremental setup (CIL), the task-id is not known during inference. Therefore we require a model to predict the task-id of the input in order to decide which set of task-specific parameters to use for inference. This prediction can be difficult, especially when there is no training data available from previous tasks. One simple approach to predict the task-id is to pass the test image via all task-specific parameters and compare the entropy of all the predictions. The prediction with the lowest entropy *i.e.*, the prediction with the peakiest confidence can be taken. However, this method does *not* work well because the cross-entropy training objective tends to make neural networks highly confident, even for out-of-distribution (OOD) samples [20]. We conjecture that such confident classifiers will predict different classes for different augmentations of an OOD data sample, but their predictions will be consistent for in-distribution samples and their augmentations. Based on this premise, the proposed approach calculates the entropy of the average predictions of a set of augmented views of the input image. The average entropy reflects the agreement of the predictions for different views of the same image, and thus a highly

confident average prediction *i.e.*, the lowest entropy average prediction, is used to infer the task identity. Once the task identity is inferred, the corresponding task-specific model is invoked on the unaugmented image. This simple trick is shown to significantly improve task prediction as verified by ablation studies (see Section 4.6). Furthermore, it does so without adding much overhead, except for the increase in batch size due to the augmentations during inference.

### 3.5. Training

As ViT is extremely data hungry and continual learning by nature is data scarce, we resorted to the compact version of ViT proposed by Hassani *et al.* [22]. In particular, we used the Compact Convolutional Transformer (CCT) which uses convolutional blocks in the tokenization step reducing the number of effective patches which in turn reduces the number of parameters used later. CCT further reduces the number of trainable parameters and improves performance by replacing class tokens with a sequence pooling strategy. Due to its effectiveness in reducing overfitting in a low data regime, we start with a CCT and train it from scratch on the initial task. Afterward, for every new incoming task, we learn (1) Task-specific convolution filters on the transformer encoder weights, (2) layernorm layers, (3) the sequence pooling layer, and (4) the final classification layer.

## 4. Experiments

We evaluated ConTraCon on several datasets in both CIL and TIL settings, comparing with many strong state-of-the-art baseline methods. We also conducted extensive ablation studies to study the effectiveness of different components in our model. Additional experimental analyses and details can be found in the supplementary material.

### 4.1. Baselines

Following [52], we compare ConTraCon against several well-established rehearsal-based Continual Learning methods such as iCARL [40], FDR [4], DER++ [6], ERT [7], RM [2]. Besides, we also compare our method with various state-of-the-art methods like Dytox [16], GPM [43], EFT [51], LvT [52], and PASS [60].

Inspired by the recent success of Class Attention in image Transformers (CAiT) [48], Dytox [16] proposed task attention-based Dynamic Token Expansion for continual classification using transformers. LvT [52] proposes a continual learning mechanism for vision transformers that utilizes an inter-task attention mechanism to consolidate knowledge from previous tasks and avoid catastrophic forgetting. Both Dytox [16] and LvT [52] include a small memory buffer to store training instances of previous tasks. Following [52], we report the performance of all the rehearsal-based approaches with memory buffer sizes of 200 and 500. GPM [43] is a regularization-based approach for TIL and

so we report only the TIL values for GPM. EFT [51] used task-specific feature-map transformation to convolutional architectures. Specifically, group-wise and point-wise convolutions were used. For a fair comparison with a similar number of parameters, we ran experiments on EFT by setting the group size of the group-wise convolution to the minimum (*i.e.*, 0) and point-wise convolution depth to 8. PASS [60] performs continual classification by learning per-class prototypes rather than storing exemplars for replay. Different augmented versions of these stored prototypes are replayed.

### 4.2. Datasets, Setup & Metrics

We evaluate the continual learning models on four benchmark datasets as follows. (1) **CIFAR-100** [27] is composed of 60K images of size  $32 \times 32$  belonging to 100 classes with 600 images per class (500 training images and the rest 100 testing images). Following [52] we divided CIFAR-100 into 5 tasks, 10 tasks and 20 tasks. (2) **TinyImageNet-200/10** [28] is a subset of the ImageNet dataset containing 100K images of size  $64 \times 64$  distributed among 200 classes. Each class has 500 training, 50 validation and 50 test images. In TinyImageNet-200/10, the 200 classes are divided into 10 sequential tasks containing 20 classes each. (3) **ImageNet-100/10** [40] contains 100 randomly chosen classes from ImageNet [42] having an average resolution of  $469 \times 387$ . It contains around 120K images for training and 5K for validation. The 100 classes are divided into 10 tasks consisting of 10 classes each. (4) **5-Datasets** [17] is composed of CIFAR-10 [27], MNIST [29], SVHN [36], Fashion MNIST [55] and notMNIST. Classification of each of these datasets is considered as a task. Additional statistics about the datasets are summarized in Table 1.

**Performance metrics:** Following [52], for Class Incremental Learning (CIL), we report top-1 accuracy over all classes of all tasks after training on the last task is completed.

For Task	Dataset	Size	Train	Test	Class
Incremental Learning (TIL), we report the average accuracy over all the tasks after training	CIFAR-100	$32 \times 32$	50K	10K	100
	TinyImageNet-200	$64 \times 64$	100K	10K	200
	ImageNet-100	$224 \times 224$	130K	5K	100
5-Datasets	CIFAR-10	$32 \times 32$	50K	10K	10
	MNIST	$32 \times 32$	60K	10K	10
	SVNH	$32 \times 32$	73K	26K	10
	FashionMNIST	$32 \times 32$	60K	10K	10
	notMNIST	$32 \times 32$	60K	10K	10

Table 1: Statistics of the benchmark datasets. on the last task. The average accuracy after training on the  $T^{th}$  task is defined by  $A_T = \frac{1}{T} \sum_{t=1}^T a_{T,t}$  where  $a_{T,t}$  is the accuracy on the test set of tasks  $t$  when the model completed learning task  $T$ .

### 4.3. Implementation Details

We used CCT [22] as our backbone architecture. For CIFAR-100, ImageNet-100 and TinyImageNet-200, we used 6 transformer encoder layers with 4 attention heads for each.

Memory Buffer	Model	Approach	Backbone	# Params	5 Tasks		10 Tasks		20 Tasks	
					TIL	CIL	TIL	CIL	TIL	CIL
200	iCARL [40]	Rehearsal	ResNet 18	11.2 M	55.70	30.12	60.81	22.38	62.17	12.62
	FDR [4]				63.75	22.84	65.88	14.84	59.13	6.70
	DER++ [6]				62.55	27.46	59.54	21.76	61.98	15.16
	ERT [7]				54.75	21.61	58.49	12.91	62.90	10.14
	RM [2]		62.05	32.23	66.28	22.71	68.21	15.15		
	LVT [52]		Transformer	8.9 M	66.92	39.68	72.80	35.41	73.41	20.63
	Dytox [16]			10.7 M	75.17	40.97	84.84	32.08	85.24	15.96
	500		iCARL [40]	Rehearsal	ResNet 18	11.2 M	64.4	35.95	71.02	30.25
FDR [4]		69.11	29.99				74.22	22.81	73.22	13.10
DER++ [6]		70.74	38.39				73.31	36.15	70.55	21.65
ERT [7]		62.85	28.82				68.26	23.00	73.50	18.42
RM [2]		69.27	39.47		73.51	32.52	75.06	23.09		
LVT [52]		Transformer	8.9 M		71.54	44.73	76.78	43.51	78.15	26.75
Dytox [16]			10.7 M		76.1	57.66	<b>88.72</b>	<b>47.34</b>	87.23	29.89
-		EFT [51]	Dynamic Arch		ResNet 18	4.9 M (32k)	79.04	<b>49.68</b>	83.14	40.42
	PASS [60]	Regulaization	ResNet 18	11.2 M	70.11	47.31	71.28	35.24	71.14	23.15
	GPM [43]	Regularization	AlexNet	6.7 M	65.90	-	72.54	-	77.59	-
	ConTraCon	Dynamic Arch	Transformer	<b>3.1 M (26k)</b>	<b>79.37</b>	48.46	85.69	41.26	<b>88.94</b>	<b>30.07</b>

Table 2: Classification accuracy on CIFAR-100 dataset. All methods except EFT, PASS, GPM and ConTraCon use a memory buffer of 200 or 500 exemplars. Accuracies in both Task Incremental Learning and Class Incremental Learning setup are reported under the columns TIL and CIL respectively. The best results are in **bold**. #Params denote the initial number of parameters in the model (in millions). For dynamic architecture based approaches the extra number of trainable parameters required per task is mentioned inside the parenthesis. Even with 500 exemplars, Dytox can outperform exemplar-free ConTraCon only on CIFAR-100/10. GPM could not be run in CIL setup and hence its accuracy is not reported.

For 5-Datasets we used 7 transformer encoder layers with 4 attention heads per layer. The encoder layers have an embedding dimension of 256. We used sinusoidal positional embeddings. For CIFAR-100 and 5-Datasets, we configured the tokenizer to a single convolution layer with kernel size 3. For TinyImageNet-200/10, the tokenizer consisted of 2 convolution layers with kernel size 5 and for ImageNet-100/10 the tokenizer had 3 convolution layers with kernel size 3. A stride of 2 and padding of 3 were used for all datasets. Following CCT [22], we set both attention and depth dropout probability to 0.1. Each task was trained for 500 epochs using AdamW [35] optimizer. We employed cosine annealing learning rate scheduler with initial learning rate of  $8e^{-4}$  and warm restarts [34]. For the convolution operation on the transformer weights, we set kernel size ( $k$ ) as 15, which is obtained by validating on a small subset of the data. Convolution is performed on the transformer weights learnt after the initial task *i.e.*,  $W_{l,h}^Q$ ,  $W_{l,h}^K$  and  $W_{l,h}^V$  in Eqn. 3 comes from the first task. For task prediction, using 10 augmentations and  $\beta = 0.6$  gives the best result.

#### 4.4. Results and Analyses

Tables 2, 3 and 4 compare ConTraCon with baselines working on both CNN and transformers, over the aforementioned four datasets. The results show the performance of our approach in both CIL and TIL setups. Table 2 shows the results on the CIFAR-100 dataset split into 5 tasks, 10 tasks, and 20 tasks following the standards commonly adopted in the community. The 5-tasks split contains fewer tasks but more classes per task. The 20-tasks split is more

challenging, where performance of most models suffer a drastic drop compared to setups consisting of fewer tasks. ConTraCon significantly outperforms existing approaches over longer sequences of tasks. In task incremental setting (TIL), ConTraCon performs the best across all the splits using the lowest amount of parameters. As rehearsal and regularization based approaches do not expand their architectures with new tasks, we provide the number of initial parameters learned for dynamic architecture expansion based approaches including ours, for fair comparison (ref # **Params** column). We also provide the number of additional parameters per task for the dynamic expansion based approaches including ours.

For the challenging CIL setup, ConTraCon does the best in the most challenging 20-Tasks scenario. For 10-Tasks, ConTraCon is the best among the dynamic architecture based approaches and better than most rehearsal based approaches. An advantage of our approach is that it does not need an extra memory buffer to store previous examples while remaining the most parameter efficient.

Table 3 shows the performance on Imagenet-100/10 and TinyImagenet-100/10. The format *dataset-X/Y* implies that the dataset contains a total of X classes divided into Y tasks uniformly. ConTraCon significantly beats all the competing approaches while remaining parameter efficient and without using replay memory. Table 4 shows the performance on the challenging 5-Datasets. While requiring more parameters compared to GPM, ConTraCon significantly outperforms it in the TIL setup. ConTraCon can equally handle CIL setup which GPM can not. In the challenging

Memory Buffer	Model	Approach	Backbone	# Params	ImageNet-100/10		TinyImageNet-200/10			
					TIL	CIL	TIL	CIL		
200	iCARL [40]	Rehearsal	ResNet 18	11.2 M	33.75	12.59	28.41	8.64		
	FDR [4]				37.80	10.08	40.15	8.77		
	DER++ [6]				31.96	11.92	40.97	11.16		
	ERT [7]				36.94	13.51	39.54	10.85		
	RM [2]				35.18	16.76	41.96	13.58		
	LVT [52]				Transformer	9.0 M	41.78	19.46	46.15	17.34
	Dytox [16]					10.7 M	70.12	41.76	61.71	19.14
500	iCARL [40]	Rehearsal	ResNet 18	11.2 M	36.89	16.44	35.89	10.69		
	FDR [4]				42.60	11.78	49.91	10.58		
	DER++ [6]				35.46	14.52	51.90	19.33		
	ERT [7]				41.56	20.42	50.87	12.13		
	RM [2]				38.66	14.56	52.08	18.96		
	LVT [52]				Transformer	9.0 M	47.84	26.32	57.93	23.97
	Dytox [16]					10.7 M	73.64	40.94	<b>64.29</b>	26.39
-	EFT [51]	Dynamic Arch	ResNet 18	4.9 M (32k)	72.18	32.98	60.00	24.08		
	PASS [60]	Regularization	ResNet 18	11.2 M	39.9	34.52	43.9	22.76		
	GPM [43]	Regularization	AlexNet	6.7 M	40.65	-	45.48	-		
	ConTraCon	Dynamic Arch	Transformer	<b>3.6 M (28k)</b>	<b>76.78</b>	<b>42.2</b>	62.76	<b>27.46</b>		

Table 3: Classification accuracy on ImageNet-100/10, TinyImageNet-200/10. All methods except EFT, PASS, GPM and ConTraCon use a memory buffer of 200 or 500 exemplars. The number of extra parameters required per task is mentioned in brackets for dynamic architecture based approaches. ConTraCon outperforms all the baselines with memory buffers 200 and 500. For ConTraCon the number of parameters for ImageNet-100/10 and TinyImageNet/10 are different. We mentioned the greater of the two in the table. GPM could not be run in CIL setup and hence its accuracy is not reported.

Model	Approach	Backbone	# Params	5-Datasets	
				TIL	CIL
Dytox [16](500)	Rehearsal	Transformer	10.7 M	77.12	<b>67.13</b>
Dytox [16](200)	Rehearsal	Transformer	10.7 M	75.81	65.04
EFT [51]	Dynamic Arch	ResNet 18	4.9 M (32k)	94.75	52.04
GPM [43]	Regularization	ResNet18	<b>1.2 M</b>	90.60	-
ConTraCon	Dynamic Arch	Transformer	3.9 M (28k)	<b>95.10</b>	65.21

Table 4: Classification accuracy on 5-Datasets. Mentioned within parentheses are the number of additional parameters required to learn each new task for dynamic architecture based approaches like EFT and ConTraCon. The memory buffer sizes for Dytox are mentioned inside the parenthesis.

CIL setup, DyTox is better but at the cost of using almost 2.5 times the number of parameters as used in ConTraCon, and using 500 exemplars (while ConTraCon is exemplar-free).

#### 4.5. Model Introspection

In this section, we analyze and explain the significance of various components of our approach. For this purpose, we use the CIFAR-100/10 dataset unless otherwise mentioned. **Upper and Lower Bounds:** We trained separate CCT backbones for each task assuming no limitations on the total number of learnable parameters. The average accuracy was found to be 89.3% which serves as the performance upper-bound having 100% increase in parameters for each new task. On the other extreme, we fine-tuned an initial CCT backbone with every new task giving 0 parameter increase per task but suffering from catastrophic forgetting the most. The average accuracy (calculated after training the last task) obtained is only 17.1%. ConTraCon achieves 85.7% aver-

age classification accuracy in TIL setup, a mere  $\sim 4\%$  drop in comparison to the upper bound while using only 0.7% of the parameters required per task by the upper bound.




Figure 4: When models are learning CIFAR-100/10 tasks in CIL setup (with 200 buffer size for exemplar-replay approaches): blue bars – the total number of parameters required (left y axis), orange bars – the average classification accuracy (right y axis). GPM could not be run in CIL setup and hence its accuracy is not reported. ConTraCon uses significantly less parameters, yet achieves higher accuracy.

**Memory Overhead:** Fig. 4 shows the number of trainable parameters required in CIFAR-100/10 dataset. As can be seen, ConTraCon has the least number of trainable parameters. For each novel task, we have separate convolution kernels, layernorm, sequence pool, and classification layers. These are cheap to store as they consist of very few parameters. ConTraCon incurs a total parameter increase of  $\sim 7\%$  for learning all the tasks in CIFAR-100/10.

**Computational Overhead:** ConTraCon uses 50% less




Figure 3: Ablation studies on ConTraCon over CIFAR-100/10 dataset: (a) Average classification accuracy using different convolutional kernel sizes. (b) Effect of the presence of local filters – interestingly, the presence of task adaptable convolution filters in the tokenizer layer lowers the performance. (c) Effect of augmentation in the inference. Reported is the average top-1 accuracy in the CIL setup calculated after training all the tasks. ConTraCon’s performance drops significantly when inference is done without augmentation. (d) Inference time vs number of augmentations for one batch of 64 images.

total FLOPs during training compared to DyTox. This is because Dytox uses additional distillation operation requiring the inputs to be passed through the transformer twice, while ConTraCon uses cheap convolution operation on transformer weights. For predicting the task during inference, we compute the entropy of average predictions of a number of augmentations of the input image. While augmentation-based inference can principally incur more computation, due to highly parallel frameworks, the overhead is less and scales linearly with number of augmentations. We ran a small experiment on CIFAR-100 with batch-size 64 with augmentations ranging from 5 to 30 in steps of 5 (Fig. 3(d)). Going from no augmentation to 10, inference time changes from 0.14s to 0.2s *i.e.*, an overhead of only 94ms per image. With the performance gain (Fig. 3(c)), our augmentation based strategy is lightweight for all practical purposes.

#### 4.6. Ablation Study

We take a closer look at our method with ablation studies on the CIFAR-100/10 dataset as described below.

**Effect of Convolution Kernel size:** We varied the kernel size from 3 to 23 (in steps of 2) and noted the average performance. Fig. 3(a) shows that the performance increases with an increase in the kernel size until the best tradeoff is reached roughly at a kernel size of 15. A further increase of the kernel size merely improves the performance by 0.2% (85.7% vis-a-vis 85.9%) but at the cost of a parameter increase of 47k per task compared to 26k with kernel size 15.

**Effect of task adaptable convolutions in Tokenizer Layer:** We wanted to see the effect of task adaptable convolution on the tokenizer layer. To this end, we divided the total number of filters in the tokenizer into global filters and local filters. The global filters are frozen after they are trained once while the local filters are trained specific to each task. the ratio of the number of global filters to total filters is termed as the *split-ratio*. A high split-ratio implies less task-specific filters whereas a low split-ratio indicates a higher number of task-specific filters and thus a higher number of task-specific parameters. A lower split ratio may adapt better to newer tasks as more parameters are dedicated to

task specific learning. However as seen in Fig. 3(b), this is not the case. We conjecture that higher number of task adaptable parameters at the stem of the network can destroy the already learnt knowledge which may be crucial at the tokenizer layer of the CCT backbone. As, highest split ratio (all global filters) gives the highest accuracy, we conducted all the experiments with the highest split ratio.

**Effect of Augmentation:** To better understand the importance of input image augmentation in task prediction, we performed a comparative analysis. As shown in Fig. 3(c), a significant drop in CIL performance can be witnessed when the task prediction is done without augmentation. This is because models trained with softmax cross-entropy loss are over-confident in their prediction [20], even if the instance does not belong to the task it is trained for. So we used average predictions of the augmentations for task prediction.

**Effect of Skip-Gating:** In order to gauge the effect of skip-gating we ran three experiments – one without any skip connection ( $\sigma(\alpha) = 0$ ), one with a learnable skip gating and one with skip connection but no gating ( $\sigma(\alpha) = 1$ ). We found the one without the skip connection got an average accuracy of 83.5%, the one with a skip connection but without gating had an average performance of 84.7% whereas the one with learnable skip-gating got an accuracy of 85.7% showing the importance of skip-gating for ConTraCon.

## 5. Conclusion

We propose Continual Transformer with Convolutions (ConTraCon), an approach that is capable of learning incoming tasks efficiently (through low-cost operations), allowing the model to expand its knowledge without blowing up its size and computation, and without having to store data samples for rehearsal. We validate the superiority of ConTraCon through extensive experimentation. We also perform thorough ablation studies highlighting the significance of various components of our approach. To our knowledge, this is one of the early works on making transformers adaptable to continual learning, and we hope our demonstration of performance and scalability will drive further research in making transformers adaptable to continual learning.



## References

- [1] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert Gate: Lifelong Learning with a Network of Experts. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3366–3375, 2017. [1](#)
- [2] Jihwan Bang, Heesu Kim, YoungJoon Yoo, Jung-Woo Ha, and Jonghyun Choi. Rainbow memory: Continual learning with a memory of diverse samples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8218–8227, June 2021. [1](#), [2](#), [5](#), [6](#), [7](#)
- [3] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The Long-Document Transformer. *arXiv preprint arXiv:2004.05150*, 2020. [2](#)
- [4] Ari Benjamin, David Rolnick, and Konrad Kording. Measuring and regularizing networks in function space. In *International Conference on Learning Representations*, 2019. [2](#), [5](#), [6](#), [7](#)
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language Models are Few-Shot Learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. [2](#)
- [6] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. In *Advances in Neural Information Processing Systems*, 2020. [2](#), [5](#), [6](#), [7](#)
- [7] Pietro Buzzega, Matteo Boschini, Angelo Porrello, and Simone Calderara. Rethinking experience replay: a bag of tricks for continual learning. *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 2180–2187, 2021. [1](#), [2](#), [5](#), [6](#), [7](#)
- [8] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end Object Detection with Transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. [1](#)
- [9] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and M Ranzato. Continual learning with tiny episodic memories. In *International Conference on Machine Learning*, 2019. [1](#), [2](#)
- [10] Krzysztof Marcin Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking Attention with Performers. In *International Conference on Learning Representations*, 2020. [2](#)
- [11] Stéphane D’Ascoli, Hugo Touvron, Matthew L Leavitt, Ari S Morcos, Giulio Biroli, and Levent Sagun. Convit: Improving vision transformers with soft convolutional inductive biases. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 2286–2296. PMLR, 18–24 Jul 2021. [2](#)
- [12] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385, 2021. [3](#)
- [13] Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Lukasz Kaiser. Universal Transformers. In *International Conference on Learning Representations*, 2018. [2](#)
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. [2](#)
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. [1](#), [2](#), [3](#)
- [16] Arthur Douillard, Alexandre Ramé, Guillaume Couairon, and Matthieu Cord. Dytox: Transformers for continual learning with dynamic token expansion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [1](#), [2](#), [5](#), [6](#), [7](#)
- [17] Sayna Ebrahimi, Franziska Meier, Roberto Calandra, Trevor Darrell, and Marcus Rohrbach. Adversarial continual learning. In *European Conference on Computer Vision*, pages 386–402. Springer, 2020. [5](#)
- [18] Mehrdad Farajtabar, Navid Azizan, Alex Mott, and Ang Li. Orthogonal gradient descent for continual learning. In *International Conference on Artificial Intelligence and Statistics*, pages 3762–3773. PMLR, 2020. [2](#)
- [19] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*, 2017. [1](#)
- [20] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR, 2017. [4](#), [8](#)
- [21] Raia Hadsell, Dushyant Rao, Andrei A Rusu, and Razvan Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in cognitive sciences*, 24(12):1028–1040, 2020. [1](#)
- [22] Ali Hassani, Steven Walton, Nikhil Shah, Abulikemu Abuduweili, Jiachen Li, and Humphrey Shi. Escaping the big data paradigm with compact transformers. *arXiv preprint arXiv:2104.05704*, 2021. [2](#), [5](#), [6](#)
- [23] Tyler L Hayes, Kushal Kafle, Robik Shrestha, Manoj Acharya, and Christopher Kanan. Remind your Neural Network to Prevent Catastrophic Forgetting. In *European Conference on Computer Vision*, pages 466–483. Springer, 2020. [1](#)
- [24] Jiangpeng He and Fengqing Zhu. Exemplar-free online continual learning. *arXiv preprint arXiv:2202.05491*, 2022. [1](#)

- [25] David Isele and Akansel Cosgun. Selective experience replay for lifelong learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. [2](#)
- [26] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. [1](#), [2](#)
- [27] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. *CiteSeer*, 2009. [5](#)
- [28] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015. [5](#)
- [29] Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010. [5](#)
- [30] Janghyeon Lee, Donggyu Joo, Hyeong Gwon Hong, and Junmo Kim. Residual continual learning. In *AAAI*, 2020. [4](#)
- [31] Sang-Woo Lee, Jin-Hwa Kim, Jaehyun Jun, Jung-Woo Ha, and Byoung-Tak Zhang. Overcoming catastrophic forgetting by incremental moment matching. *Advances in neural information processing systems*, 30, 2017. [2](#)
- [32] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. [1](#)
- [33] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient Episodic Memory for Continual Learning. *Advances in neural information processing systems*, 30, 2017. [1](#)
- [34] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. [6](#)
- [35] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. [6](#)
- [36] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011. [5](#)
- [37] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019. [1](#)
- [38] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. GDumb: A Simple Approach that Questions our Progress in Continual Learning. In *European conference on computer vision*, pages 524–540. Springer, 2020. [1](#)
- [39] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? In *Advances in Neural Information Processing Systems*, 2021. [4](#)
- [40] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017. [1](#), [2](#), [5](#), [6](#), [7](#)
- [41] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019. [1](#)
- [42] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3), 2015. [5](#)
- [43] Gobinda Saha, Isha Garg, and Kaushik Roy. Gradient projection memory for continual learning. In *International Conference on Learning Representations*, 2021. [2](#), [5](#), [6](#), [7](#)
- [44] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, pages 4548–4557. PMLR, 2018. [1](#)
- [45] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 2994–3003, 2017. [2](#)
- [46] Pravendra Singh, Vinay Kumar Verma, Pratik Mazumder, Lawrence Carin, and Piyush Rai. Calibrating cnns for lifelong learning. *Advances in Neural Information Processing Systems*, 33:15579–15590, 2020. [1](#)
- [47] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers and distillation through attention. In *International Conference on Machine Learning*, volume 139, pages 10347–10357, July 2021. [2](#)
- [48] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 32–42, 2021. [5](#)
- [49] Gido M. van de Ven and Andreas S. Tolias. Three scenarios for continual learning. *CoRR*, abs/1904.07734, 2019. [3](#)
- [50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. [2](#)
- [51] Vinay Kumar Verma, Kevin J Liang, Nikhil Mehta, Piyush Rai, and Lawrence Carin. Efficient feature transformations for discriminative and generative continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13865–13875, 2021. [2](#), [5](#), [6](#), [7](#)
- [52] Zhen Wang, Liu Liu, Yiqun Duan, Yajing Kong, and Dacheng Tao. Continual learning with lifelong vision transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 171–181, 2022. [1](#), [2](#), [5](#), [6](#), [7](#)
- [53] Zifeng Wang, Zizhao Zhang, Sayna Ebrahimi, Ruoxi Sun, Han Zhang, Chen-Yu Lee, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, et al. Dualprompt: Complementary

- prompting for rehearsal-free continual learning. *European Conference on Computer Vision*, 2022. 3
- [54] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 139–149, 2022. 3
- [55] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv*, 2017. 5
- [56] Mengqi Xue, Haofei Zhang, Jie Song, and Mingli Song. Meta-attention for vit-backed continual learning. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [57] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [58] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. In *International Conference on Learning Representations*. ICLR, 2018. 1, 2
- [59] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR, 2017. 2
- [60] Fei Zhu, Xu-Yao Zhang, Chuang Wang, Fei Yin, and Cheng-Lin Liu. Prototype augmentation and self-supervision for incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5871–5880, June 2021. 5, 6, 7